

Lessons Learned from a Ligand-Unbinding Stress Test for Weighted Ensemble Simulations

Anthony T. Bogetti, Darian T. Yang, Hannah E. Piston, David N. LeBard, and Lillian T. Chong*



Cite This: <https://doi.org/10.1021/acsomega.5c03809>




Read Online

ACCESS |

Metrics & More

Article Recommendations

ABSTRACT: The weighted ensemble (WE) path sampling strategy has pushed the boundaries of molecular simulation by enabling the generation of rates and atomistic pathways for biological processes beyond the ms time scale. However, the WE strategy has not yet reached its full potential and much can be gained from pursuing “stress tests”. Here, we have explored a stress test involving the seconds-timescale unbinding of a highly charged ligand from a protein receptor: the release of the ADP ligand from the Eg5 protein receptor, which functions as a motor protein in cell division. From this stress test, we learned valuable lessons regarding the choice of progress coordinate and improvements to the WE resampling procedure. Based on the latter, we have developed a WE method referred to as the minimal adaptive binless (MABL) method. The MABL method is in the same spirit as our previously developed minimal adaptive binning scheme for surmounting large energetic barriers but is “binless”, i.e., does not require the use of rectilinear bins along a progress coordinate. This minimal version of a binless method is >50% more efficient than the corresponding binned version and provides a framework for implementing more complex binless methods.



INTRODUCTION

The weighted ensemble (WE) path sampling strategy^{1,2} has enabled the generation of pathways with rigorous kinetics for grand-challenge applications in molecular simulation. These applications include large-scale conformational transitions within proteins,^{3,4} protein folding,^{5–8} protein–protein binding,^{9,10} protein–ligand unbinding,^{11,12} and chemical reactions.¹³ The WE strategy efficiently samples barrier-crossing processes by running multiple, weighted trajectories in parallel and periodically applying a resampling procedure to provide even coverage of configurational space.^{1,2} Typically, configurational space is divided into regions, or “bins” along a progress coordinate—such as the distance between residues or the RMSD to a target structure—that captures the system’s slow, relevant motions. Although progress coordinates can be multidimensional, those with more than three dimensions are often impractical due to computational limitations.

To address this challenge, “binless” WE methods have been developed. These approaches reduce the cost of tracking progress in high-dimensional spaces by using a one-dimensional scoring functions instead of multidimensional progress coordinates.^{7,8,14,15} While binless WE methods still rely on metrics of progress, they apply these metrics in a more efficient way compared to binned WE techniques. A key feature of the WE strategy is that the progress coordinate can be changed “on the

fly” during a simulation, as trajectory weights are independent of the specific progress coordinate used.¹⁶

Enhanced sampling strategies generally aim to make molecular simulations more efficient by reducing the time spent exploring redundant stable states. These strategies promote broader sampling through one of three main approaches. The first class of strategies uses elevated temperatures to help systems overcome local barriers, with parallel tempering strategies—such as replica exchange molecular dynamics—being a widely used example. The second class introduces biasing forces to help the system surmount barriers, as exemplified by metadynamics. The third class, which includes WE and other path sampling strategies, enhances sampling of transitions between stable states rather than between the states themselves without altering the free energy landscape or applying any biasing forces. Strategies in the first class cannot provide rate estimates. Methods of the second class—such as τ -random acceleration molecular dynamics (τ -RAMD),¹⁷ scaled-

Received: April 25, 2025

Revised: June 3, 2025

Accepted: June 4, 2025

MD,¹⁸ targeted MD,¹⁹ and steered MD²⁰—typically yield only relative rate estimates. Some approaches within the second and third classes, particularly metadynamics,²¹ Gaussian accelerated MD,²² and path sampling strategies more generally, can provide absolute rate estimates. The WE strategy is unique among these methods in its ability to provide absolute rate estimates while remaining highly flexible.¹¹ Unlike post-simulation methods such as Markov state modeling or path sampling strategies such as weighted ensemble milestone (WEM) that prioritize rate estimates with discontinuous trajectory segments,^{23,24} the WE strategy emphasizes the generation of continuous pathways. These detailed pathways not only yield rate estimates but also enable rich mechanistic insights.

Like any enhanced sampling method, there is no “free lunch” for WE path sampling over the conventional “brute-force” manner of running sufficiently long molecular dynamics (MD) simulations to capture the process of interest. The main caveat of the WE strategy is that key motions of the process of interest may be orthogonal to the progress coordinate and sampled in a brute-force manner. That said, effective progress coordinates have been identified for a variety of complex biological processes on the seconds-time scale or beyond. For example, WE simulations have generated pathways and rates for a process involving the escape of an uncharged drug-like ligand from a completely buried cavity of a protein receptor.¹¹ These simulations involved a combination of progress coordinates monitoring (i) the opening of the receptor cavity using the solvent-accessible surface area of the cavity, (ii) the relative orientation of the ligand and receptor using an “unbinding” root mean squared deviation (RMSD) involving the heavy-atom RMSD of the ligand after alignment on the receptor in the bound state, and (iii) the ligand–receptor distance using the minimum distance between the two binding partners. To efficiently surmount barriers, bins were adaptively positioned along these coordinates using the minimal adaptive binning (MAB) method.²⁵

Here, we present lessons learned from a ligand-unbinding application in which the WE protocol described above for uncharged ligands failed to generate successful events for a charged ligand. Our application was a “stress test” for the WE strategy and involved the unbinding of a highly charged ligand from a protein receptor, i.e., the release of an ADP ligand from the Eg5 motor protein (Figure 1), which is the rate-limiting step in a key process of cell division.²⁸ The generation of fully atomic, continuous ligand-unbinding pathways and the prediction of ligand k_{off} values have long been of interest to the drug discovery pipeline given that drug efficacy is often correlated with the inverse of the k_{off} , i.e., residency time of a ligand in the receptor binding site.²⁹

The simulation of ADP unbinding serves as a relevant and challenging model system for developing methods to predict k_{off} values. Our lessons learned provide insights into the choice of progress coordinate, manner of applying adaptive binning, and the WE resampling procedure. Based on the latter lesson, we have developed a minimal adaptive binless WE method (called MABL) that generalizes our MAB scheme.²⁵ Our MABL method enables the efficient generation of continuous pathways for our highly charged ligand-unbinding stress test. The MABL method framework, along with many of the MABL variants tested for this project, can be found at the following GitHub repository: https://github.com/westpa/user_submitted_scripts/tree/main/MABL.




Figure 1. ADP-bound Eg5 motor protein. Crystal structure of the ADP-bound Eg5 motor protein (PDB code: 1II6).²⁶ As a rate-limiting step for cell division, ADP release from Eg5 has been targeted for the design of allosteric inhibitors.²⁷

THEORY

In this section, we provide a brief overview of the WE strategy and outline the development of our minimal adaptive binless WE resampler, the MABL method, which allowed us to sample unbinding pathways for the highly charged ADP ligand. The MABL method is an extension of the MAB strategy for binless resampling, which provides greater flexibility and efficiency gains compared to the MAB method.

Weighted Ensemble (WE) Strategy. The WE strategy involves running multiple weighted trajectories in parallel and periodic application of a resampling procedure at fixed short time intervals τ . This procedure involves replicating and terminating trajectories to provide an even coverage of configurational space. Trajectories that populate less-visited regions of configurational space are replicated to enhance their success, dividing their statistical weight among the “split” trajectories. Trajectories that populate already visited regions are terminated, merging their statistical weight with that of another trajectory in the same region. The combination of running MD for a resampling time interval of τ and applying the resampling procedure constitutes a single WE iteration.

Minimal Adaptive Binning (MAB) Method with Multiple Regions. To efficiently surmount high energy barriers, we previously developed the minimal adaptive binning (MAB) method for adaptive placement of bins along a progress coordinate during a WE simulation.²⁵ This method populates less-visited regions of state space by adaptively placing bins near the leading edge of progress and near bottleneck regions. Bins are also evenly spaced between the lagging and leading edges of the progress coordinate. Bin positions are determined after each resampling time interval τ .

For systems with particularly large barriers, one issue with the MAB method is the generation of trajectories with extremely low weights due to oversplitting—that is, repeatedly splitting the same leading trajectories. These very low weights can hinder convergence to a steady state, which is essential for obtaining reliable rate estimates. To reduce the likelihood of oversplitting, a multi-MAB scheme can be employed.^{11,30} In this approach, the

progress coordinate is divided into several, relatively large regions, and a separate MAB scheme is nested within each region (Figure 2). This use of multiple, independent MAB




Figure 2. Basic illustration of the multi-MAB method. In the top schematic, a single MAB scheme (MAB method) covers the entire span of the progress coordinate, which can yield trajectories with extremely low weights due to oversplitting (i.e., repeated splitting of the same) leading trajectories. In the bottom scheme, the progress coordinate is divided into multiple regions, and a separate MAB scheme is nested in each region (multi-MAB method). Nesting multiple MAB schemes in larger regions reduces the generation of leading trajectories with extremely low weights.

schemes is conceptually similar to other path sampling strategies, such as transition interface path sampling³¹ and the use of WE simulations between milestones in the weighted ensemble milestones (WEM) method.³² The larger regions in which MAB bins are nested are defined by the user prior to the simulation and should reflect meaningful physical aspects of how the coordinates intersect. For example, in a recent WE application to the escape of an uncharged, drug-like ligand from a completely buried protein receptor,¹¹ the receptor cavity solvent-accessible surface area (SASA) needed to be binned more finely when unbinding RMSD was low because at high unbinding RMSD, the cavity would have had to already be open. Therefore, one MAB region was placed at a low unbinding RMSD and another at a high unbinding RMSD.

Minimal Adaptive Binless (MABL) Method. While the MAB method is efficient in surmounting large barriers, the computational expense of adaptive binning greatly increases with multidimensional progress coordinates. To enable efficient WE resampling of the high-dimensional state space, we have developed a minimal adaptive binless (MABL) method for WE simulations. In contrast to binless WE methods such as the Resampling of Ensembles by Variance Optimization (REVO) method,¹⁴ which utilizes pairwise RMSD “distances” to split and merge trajectories, our MABL method is a minimal, intuitive binless scheme. Similar to our binned MAB method, the MABL method splits trajectories at the leading edge along one or multiple progress coordinates.

At the heart of the MABL method is the use of a progress score, denoted as S , which quantifies the combined progress along multiple coordinates at a computational cost comparable to tracking a single coordinate. This score is calculated for each trajectory i after propagating dynamics for a short time interval τ and reflects the trajectory’s advancement—on a scale of 0 to 1—along each m coordinate of progress q between an initial value $q_{m,i}$ and target value $q_{m,t}$ along M total coordinates:

$$S_i = \prod_{m=1}^M C_m \left(1 - \left| \frac{q_m - q_{m,t}}{q_{m,i} - q_{m,t}} \right| \right) \left(\frac{1}{-\ln(P_i)} \right) \quad (1)$$

where q_m is the current value of coordinate m of trajectory i . The initial and target values of each coordinate are user-defined at the start of the simulation and can be adjusted on the fly during a simulation. The score S is calculated as a product of the progress along each individual coordinate, enabling the inclusion of multiple coordinates in a single metric. This allows for flexible and comprehensive monitoring of the overall simulation progress. The scaling factors C_m enable users to tune the influence of each coordinate on the overall progress score. This adjustment serves a similar purpose to the multi-MAB scheme: to emphasize or de-emphasize certain progress measures

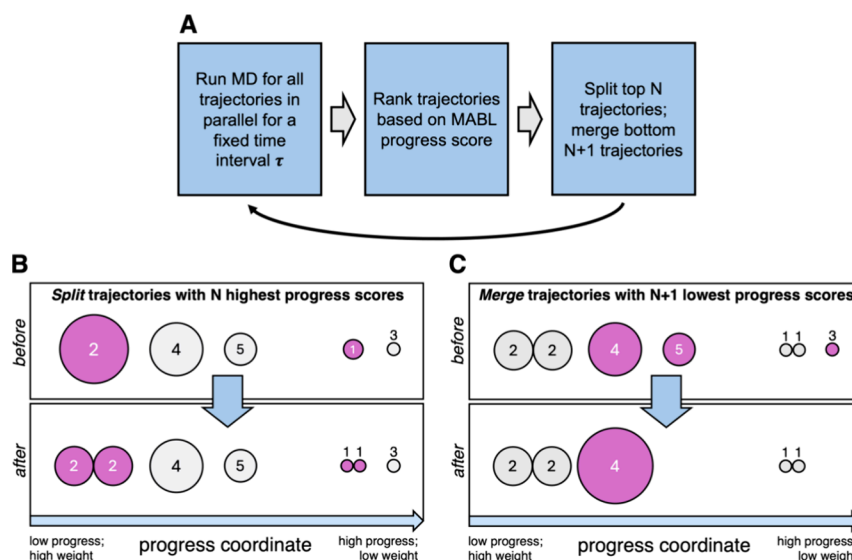


Figure 3. The minimal adaptive binless (MABL) method for WE simulations. (A) Workflow for the MABL method. (B) Illustration of the splitting step in the WE resampling procedure. The “before” box shows a trajectory distribution after the resampling time interval τ , where trajectories with varying weights (represented by circle size) are ranked by a progress score S (see eq 1; in this example, the only coordinate in S is the distance along the x axis). The top $N = 2$ trajectories (trajectories #1 and #2) are split, yielding the child trajectories in the “after” box. (C) Illustration of the merging step in the WE resampling procedure. To maintain a fixed total number of trajectories (here, five trajectories), the bottom $N + 1 = 3$ trajectories are merged.

depending on the expected relevance at difference stages of the simulated process.

To reduce the likelihood of oversplitting leading trajectories, a balance term $1/(-\ln(P))$ is included, where P is the statistical weight of the trajectory being considered. By taking the fractional, negative log of P , lower-weight trajectories would be less favored for splitting and higher-weight trajectories would be more favored for splitting, thereby facilitating the flow of probability away from the initial state.

As illustrated in Figure 3, the WE resampling procedure in the MABL method begins by ranking all trajectories by their progress score S . To maintain a constant number of trajectories, the top N scoring trajectories are split, while the bottom N scoring trajectories are merged. Although the optimal value of N was not extensively benchmarked in our initial simulations, we found that $N = 5$ consistently facilitated the generation of productive ADP unbinding pathways. Based on this observation, we recommend choosing N to be ~ 10 – 20% of the total number of trajectories in the WE simulation. It is important to note that N cannot exceed 50% of the total number of trajectories as a single trajectory cannot be selected for both splitting and merging in the same resampling step. In accordance with WE resampling rules, when trajectories are merged, the surviving trajectory is selected probabilistically by the statistical weights of the candidates being considered.

Compared with a binned WE method, our binless MABL method greatly reduces the number of adjustable parameters for running a WE simulation, especially for more complex processes such as receptor–ligand unbinding. In particular, the use of a single progress score involves (i) defining only a range of values for the progress coordinate rather than dividing a progress coordinate into bins and, optionally, (ii) specifying the relative importance of different progress coordinates rather than the complex task of nesting a progress coordinate within another progress coordinate.

METHODS

Preparation of the Initial State Ensemble. All system preparation was performed using the tleap program of the Amber software package.³³ The Eg5-ADP complex was prepared starting from heavy-atom coordinates of the complex from the crystal structure (PDB: 1II6),²⁶ protonating titratable residues for neutral pH, and then solvating the system in a truncated octahedral box of explicit water molecules with a 16 Å clearance between the complex and the edge of the box. To mimic the experimental salt concentration²⁸ of 100 mM NaCl, 73 Na⁺ ions and 73 Cl[−] ions were placed in the solvent box. The protein was treated with the Amber ff19SB force field³⁴ and the waters with the OPC water model.³⁵ The ions, including a single Mg²⁺ ion in complex with ADP, were treated with Li-Merz 12–6–4 parameters compatible with OPC water,³⁶ and the ADP ligand was treated with parameters compatible with the Amber ff19SB force field.³⁷ Short-range interactions were truncated at 10 Å, and long-range electrostatic interactions were treated using the particle mesh Ewald method.³⁸ The system was energy-minimized for 10,000 steps before heating for 20 ps with position restraints on all solute heavy atoms to 298 K at constant volume with a weak Langevin thermostat using a collision frequency of 1 ps^{−1}. Following heating, the system was equilibrated for 1 ns with position restraints on all solute heavy atoms at a constant pressure with a Monte Carlo barostat, with pressure changes attempted every 0.2 ps. Finally, an

additional 1 ns of constant-pressure equilibration was performed after removal of the solute heavy-atom restraints.

WE Simulations. All WE simulations were run using the WESTPA 2.0 software package,³⁹ a resampling time interval of 50 ps, and initiated from the equilibrated bound structure of the Eg5/ADP system. All MD simulations were run using the pmemd.cuda GPU-accelerated dynamics engine of the Amber 22 package³³ at a constant temperature (298 K) and pressure (1 atm) using the weak Langevin thermostat and Monte Carlo barostat mentioned above for the equilibration procedure.

The multidimensional progress coordinate used in the MABL method incorporated the following metrics, each representing a different dimension of progress.

- **Unbinding RMSD.** The heavy-atom RMSD of the ligand, calculated after aligning the protein receptor to the initial bound-state structure. Alignment was restricted to receptor regions with relative low RMS fluctuations (residues 18–364), as identified by conventional MD simulations.
- **Ligand–receptor interaction energy (E_{int}).** The non-bonded interaction energy between the ligand and the receptor, including both van der Waals and electrostatic contributions. This was defined as the total energy of the ligand–receptor complex minus the sum of the individual energies of the ligand and receptor.
- **Ligand–receptor distance.** The minimum separation distance between any atom of the receptor and ligand.

This last metric was primarily used to define key states such as the unbound state, bound state, and, if necessary, an encounter-complex intermediate. While it may have been conceptually cleaner to test each coordinate individually within the MABL framework, all three coordinates were ultimately necessary to capture the ADP unbinding process.

We compare our MABL WE simulation with a WE simulation run with a multi-MAB strategy. The multi-MAB WE simulation was run with a target number of four trajectories per bin. A four-dimensional progress coordinate was employed consisting of (i) the unbinding RMSD of the ADP ligand after alignment on the Eg5 receptor, (ii) the ligand–receptor interaction energy, (iii) the ligand–receptor separation distance, and (iv) the minimum separation distance between the phosphate tail of the ADP ligand and the Eg5 receptor. For each dimension of this progress coordinate, outer bin boundaries of [0, 7.5, 10.5, 'inf'], ['−inf', 10 'inf'], [0, 6, 'inf'], and [0, 6, 'inf'] were defined.

Four MAB schemes were nested within the outer bins. The first MAB scheme was intended to sample increasing ligand RMSD and was placed at [3, 55, 5, 5] and contained [5, 1, 1, 1] MAB bins per dimension with directions of [1, −1, 1, 1]. The second MAB scheme was intended to sample increasing ligand–receptor distance to higher RMSD values and was placed at [8, 55, 5, 5] and contained [5, 1, 5, 1] MAB bins per dimension with directions of [1, −1, 1, 1]. The third MAB scheme was intended to sample interaction energies after the RMSD had been increased over 10 Å and was placed at [11, 55, 5, 5] and contained [1, 5, 1, 1] MAB bins per dimension with directions of [1, −1, 1, 1]. The fourth and final MAB scheme was intended to focus sampling on increasing the minimum distance between the ligand's phosphate tail and the receptor while also still focusing on sampling the interaction energy and was placed at [11, 5, 5, 5] and contained [1, 5, 1, 5] MAB bins per dimension with directions of [1, −1, 1, 1].

The MABL WE simulation was conducted with a fixed number of 40 trajectories per WE iteration. Progress scores in MABL were calculated using three quantities: (1) the RMSD of the ligand following alignment on the receptor, (2) the interaction energy between the ligand and receptor, and (3) the minimum distance between the ligand and receptor. In contrast to the multi-MAB simulations, the minimum distance between the ligand's phosphate tail and the receptor was excluded from the MABL simulations. This choice was based on extensive testing, which revealed that the measurement was not necessary to achieve successful unbinding events. Since the metric was ultimately not needed, we do not expect it to have significantly impacted sampling efficiency in the multi-MAB simulations, and thus, comparisons between MABL and multi-MAB simulations remain valid. For the RMSD, progress was evaluated between 0 and 25 Å; for the interaction energy, between 350 and −200 kcal/mol; and for the ligand–receptor distance, between 0 and 10 Å. In addition, to prevent the ligand from becoming “trapped” in nearby regions of the binding pocket, trajectories with RMSD values between 10 and 13 Å had their progress scaled by a factor of 0.8. These bounding values were selected based on preliminary WE simulations using an MAB scheme or multi-MAB scheme.

RESULTS AND DISCUSSION

In this section, we present in detail the lessons learned from our successes and failures from WE simulations of a stress test: the unbinding process involving the charged ADP ligand from the Eg5 protein receptor. Among 44 different protocols for WE simulations of ADP unbinding from Eg5, we generated over 43.8 μ s of aggregate simulation time. In this set of WE simulations, we tested a variety of progress coordinates, binning schemes, and criteria for the WE resampling procedure. We learned two main lessons for applying WE simulations to study ligand unbinding that are particularly relevant to charged ligands. These lessons led to the development of our binless WE method, the MABL method, which was instrumental in enabling the efficient generation of ADP unbinding events.

Lesson #1: Both Energetic and Structural Features Are Necessary for Monitoring Unbinding of a Charged Ligand. Effective Progress Coordinates for the Unbinding of Uncharged Ligands Are Not Sufficient for Charged Ligands. While previous WE studies have found the unbinding RMSD to be an effective progress coordinate for simulating unbinding events with uncharged ligands,¹¹ our tests revealed that this metric was not an effective progress coordinate for the unbinding of the charged ADP ligand. This result is due to the fact that the unbinding RMSD does not provide a continuous range of values that is sufficiently large to capture incremental amounts of progress toward the target unbound state. The detection of incremental amounts of progress is crucial for charged ligands given their tighter binding affinity for the protein receptor.

The WE strategy, when using a multidimensional progress coordinate, is most effective when one dimension has many “units” of progress toward the target state that simultaneously contribute to larger units of progress in other dimensions. For instance, in the case of receptor–ligand unbinding, increases in ligand–receptor interaction energy should (at least eventually) translate to increases in unbinding RMSD, which, in turn, translate to increases in receptor–ligand distance (the main, and most broad, determinant of success). In cases where the ligand is uncharged, unbinding RMSD may be sufficient as a progress

coordinate because progress in unbinding RMSD can directly contribute to progress in receptor–ligand distance. However, in the case of charged-ligand unbinding, the tighter interactions between the ligand and receptor compress the range of progress “units” available to the unbinding RMSD, rendering it not as effective. To overcome this, we identified and employed a more relevant coordinate with an order of magnitude more progress units that the WE strategy could utilize: the interaction energy.

Interaction Energy Is an Effective Progress Coordinate for Sampling Charged-Ligand Unbinding. The ligand–receptor interaction energy, with a large range of possible values due to electrostatic interactions, was effective as a progress coordinate for charged-ligand unbinding. By considering the interaction energy along with the unbinding RMSD and the ligand–receptor distance, we were able to get close to generating ligand-unbinding pathways with the MAB method but were heavily limited by the large computational cost of binning in multiple dimensions of the progress coordinate. With the inclusion of the interaction energy as a coordinate in the MABL progress score, we were finally able to generate ligand-unbinding events (Figure 4).

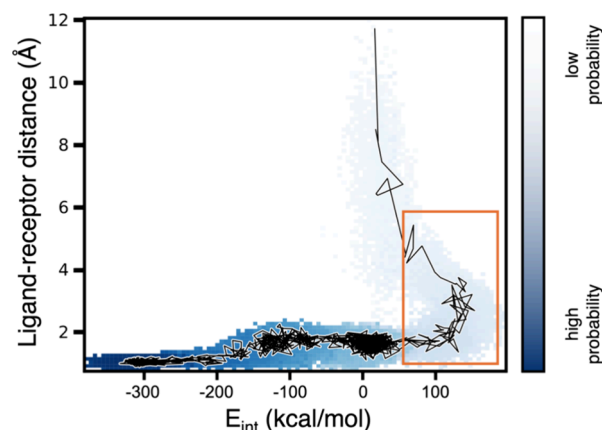


Figure 4. Repulsive ligand–receptor interactions facilitate ligand unbinding. Probability distribution as a function of the ligand–receptor interaction energy (E_{int}) and ligand–receptor separation distance for a single WE simulation of the Eg5-ADP unbinding process. A representative unbinding pathway is traced in black. Repulsive ligand–receptor interactions ($E_{\text{int}} > 0$; delineated in red) appear to facilitate the dissociation of the ADP ligand from the Eg5 protein receptor. Given the extremely low probabilities in this simulation, which is far from a steady state, we interpret this distribution qualitatively.

Backward Progress Can Be Essential. As is evident in Figure 4, repulsive interactions between the ligand and receptor occur immediately before the ligand dissociates from the receptor. Such interactions may serve as a “springboard” for launching the ligand away from the charged regions of the receptor binding pocket. Said another way, backward motion along a progress coordinate can be not only useful, but essential for enhancing the sampling of pathways for unbinding of the charged ADP ligand.

In our initial efforts with the MABL method that employed a progress score consisting of the interaction energy, unbinding RMSD, and ligand–receptor distance, our WE simulations resulted in trapping of the ligand–receptor system in conformations with large unbinding RMSD values. These trapped conformations corresponded to potential encounter

complexes in which the ligand was bound to the receptor but not in its native bound pose.

To enable successful unbinding events, we allowed backward progress along the unbinding RMSD coordinate by reducing its relative contribution to the overall progress score by 20% in the range of 10 to 13 Å, where the ligand was getting trapped in a non-native bound conformation. After implementing this adjustment, the ADP ligand was able to fully dissociate from the Eg5 receptor, as shown in Figure 4.

Lesson #2: While Our Binless WE Method Is More Efficient in Capturing Rare Events, Care Must Be Exercised in the Merging of Trajectories. Our Minimal Adaptive Binless (MABL) Method Enables Efficient WE Simulations with a High-Dimensional Progress Coordinate. Using a progress score/coordinate that consists of the ligand–receptor interaction energy, unbinding RMSD, and the ligand–receptor distance, our binless MABL method was >50% more efficient at generating the first ligand-unbinding event compared to our most effective binned resampler employing the multi-MAB scheme. The aggregate simulation time needed for the MABL method to generate the first ligand-unbinding event was only 2.0 μ s compared with 4.7 μ s for the multi-MAB method. In addition, our MABL method maintains a fixed number of trajectories throughout the simulation—a desirable feature that facilitates the planning of computational resources to allocate for a WE simulation.

Care Must Be Taken when Merging Trajectories within a Binless Framework. While the progress score used in the MABL method is primarily focused on the splitting of promising trajectories, care must be exercised in the groupings of trajectories to consider for merging within a binless WE method for several reasons.

First, if nonredundant trajectories are grouped for merging, trajectories that are very different from the surviving trajectory could be terminated. The loss of nonredundant trajectories could hamper sampling of “breakout events” and generate a less-diverse path ensemble. This issue has been addressed for the original Huber and Kim WE framework in the recent development of an equal-weight resampler.⁴⁰

Second, the total probability may become concentrated in just one or a few trajectories, leading to much lower-than-average probabilities of trajectories, resulting in much lower-than-average probabilities for those at the leading edge of sampling. We were able to alleviate—but not entirely eliminate—this accumulation by incorporating trajectory weights into the progress score. Additional strategies to more effectively prevent large-scale probability accumulation will be explored in future work.

Based on our efforts to refine the WE resampling procedure in our MABL method, we present a cautionary example where we violated one of the two rules that must be followed by a resampling procedure to avoid introducing statistical bias into the WE simulation:

1. Trajectory weights must always sum to a total probability of one, and
2. When evaluating trajectories for merging, the surviving trajectory must be chosen according to its weight (probability) or randomly chosen, if more than one candidate exists with the same weight.¹

In our example, we broke the second rule when attempting to select the surviving trajectory based on their progress score. While the progress score contains a weight-based balance term

(see Methods), this criteria for selecting the surviving trajectory resulted in the termination of a higher-weight trajectory and the subsequent merging of its weight onto a lower-weight surviving trajectory. This merging event led to an unphysical result where there was no reduction in weights as the trajectories approached the target state (i.e., as if no barriers exist in the unbinding process) due to two high-weight trajectories ending up on a “fast track” from the initial to the target state (Figure 5). Our example

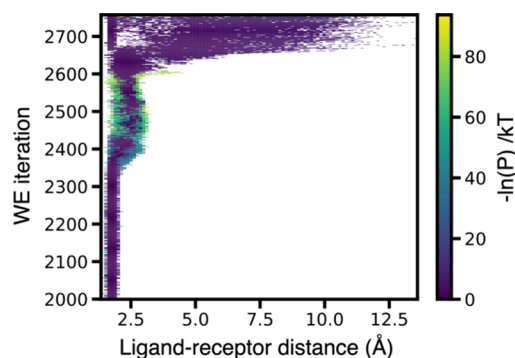


Figure 5. Time-evolution of the probability distribution as a function of ligand–receptor distance from a WE simulation in which the choice of trajectories for merging breaks the rules. The preferential merging of heavy-weight trajectories in an initial version of the MABL method violated the statistical rules of the WE resampling procedure.¹ As a result, a few heavy-weight trajectories make a “bee-line” from the initial to target state without any splitting of the trajectories.

underscores the importance of following not just the first rule but also the second rule for the WE resampling procedure that involves the random selection of trajectories that survive a merging event.

Strategy for Enforcing “Soft” Threshold Values for Trajectory Weights. In certain WE studies, the use of minimum and maximum threshold values for trajectory weights has prevented the weights from becoming extremely small and dominating the path ensemble, respectively.^{12,14} However, hard limits on trajectory weights can hinder progress toward the target state by preventing any splitting of “breakaway” trajectories. To achieve a middle ground between being too strict and too relaxed with trajectory weights, we introduced “soft” thresholds into the MABL method by incorporating trajectory weights into the progress score. Based on this progress score, higher-weight trajectories are chosen for splitting, and lower-weight trajectories are considered for merging. The “soft” threshold term used in our progress score can be further modified and serves as a starting point for exploring the use of this more flexible form of trajectory weight thresholds.

CONCLUSIONS

We have presented lessons learned from WE simulations of a challenging stress test: the unbinding of a charged ligand, specifically the release of ADP from the Eg5 motor protein. Our lessons are drawn from >43.8 μ s of aggregate WE simulation time across 44 different simulation protocols. Many of our unsuccessful attempts ultimately led to the development of a new WE strategy—the MABL method—a binless version of our previously developed MAB method.

The MABL method performs WE resampling using a single progress score that integrates progress along multiple coordinates. This binless approach proved to be more efficient than our best binned strategies when handling multiple progress

coordinates. Further, MABL maintains a fixed number of total trajectories throughout a simulation, which simplifies computational resource allocation.

Despite the complexity of simulating the unbinding of a charged ligand, we successfully generated unbinding pathways using the MABL method by incorporating three key features: (i) inclusion of ligand–receptor interaction energy in the progress score, (ii) enabling backward progress by reducing the weight of the unbinding RMSD coordinate where the ligand was prone to becoming trapped, and (iii) mitigating oversplitting at the leading edge of sampling by integrating trajectory weights into the progress score. Although rate constants could not yet be estimated from these initial pathways, they represent an important first step toward that goal.

Finally, we implemented our MABL method within the open-source WESTPA software package, providing a framework for future development of more advanced binless strategies to address complex biomolecular processes.

AUTHOR INFORMATION

Corresponding Author

Lillian T. Chong – Department of Chemistry, University of Pittsburgh, Pittsburgh, Pennsylvania 15260, United States; orcid.org/0000-0002-0590-483X; Email: ltchong@pitt.edu

Authors

Anthony T. Bogetti – Department of Chemistry, University of Pittsburgh, Pittsburgh, Pennsylvania 15260, United States; Present Address: Laufer Center for Physical and Quantitative Biology, Stony Brook University, Stony Brook, NY 11794, United States; orcid.org/0000-0003-0610-2879

Darian T. Yang – Department of Chemistry, University of Pittsburgh, Pittsburgh, Pennsylvania 15260, United States; Present Address: Department of Biology Biocenter, Room 1.0.20 Ole Maaløes Vej 5 DK-2200 Copenhagen N, Denmark; orcid.org/0000-0002-8654-3529

Hannah E. Piston – Department of Chemistry, University of Pittsburgh, Pittsburgh, Pennsylvania 15260, United States

David N. LeBard – OpenEye, Cadence Molecular Sciences, Santa Fe, New Mexico 87508, United States; orcid.org/0000-0002-0979-800X

Complete contact information is available at: <https://pubs.acs.org/10.1021/acsomega.5c03809>

Author Contributions

The manuscript was written through contributions of all authors. All authors have given approval to the final version of the manuscript.

Funding

This work was supported by NIH grant R01 GM1151805 to L.T.C. and University of Pittsburgh Andrew Mellon Predoctoral Fellowships to A.T.B. and D.T.Y. Support was also provided by an OE Scientific Graduate Student Internship involving A.T.B. under the mentorship of D.N.L. Computational resources were provided by NSF XSEDE allocation TG-MCB100109 to L.T.C. and the University of Pittsburgh Center for Research Computing, RRID:SCR_022735, through the H2P cluster, which is supported by NSF award number OAC-2117681.

Notes

The authors declare the following competing financial interest(s): LTC is a member of the Scientific Advisory Board for OpenEye Scientific, Cadence.

ACKNOWLEDGMENTS

We thank AJ Pratt for initial efforts and Anthony Nicholls for insightful discussions.

REFERENCES

- (1) Huber, G. A.; Kim, S. Weighted-ensemble Brownian dynamics simulations for protein association reactions. *Biophys. J.* **1996**, *70*, 97–110.
- (2) Zuckerman, D. M.; Chong, L. T. Weighted Ensemble Simulation: Review of Methodology, Applications, and Software. *Annu. Rev. Biophys.* **2017**, *46*, 43–57.
- (3) Bogetti, X.; Bogetti, A.; Casto, J.; Rule, G.; Chong, L.; Saxena, S. Direct observation of negative cooperativity in a detoxification enzyme at the atomic level by Electron Paramagnetic Resonance spectroscopy and simulation. *Protein Sci.* **2023**, *32*, No. e4770.
- (4) Sztain, T.; Ahn, S. H.; Bogetti, A. T.; Casalino, L.; Goldsmith, J. A.; Seitz, E.; McCool, R. S.; Kearns, F. L.; Acosta-Reyes, F.; Maji, S.; et al. A glycan gate controls opening of the SARS-CoV-2 spike protein. *Nat. Chem.* **2021**, *13*, 963–968.
- (5) Adhikari, U.; Mostofian, B.; Copperman, J.; Subramanian, S. R.; Petersen, A. A.; Zuckerman, D. M. Computational Estimation of Microsecond to Second Atomistic Folding Times. *J. Am. Chem. Soc.* **2019**, *141*, 6519–6526.
- (6) Santhouse, J. R.; Leung, J. M. G.; Chong, L. T.; Horne, W. S. Implications of the unfolded state in the folding energetics of heterogeneous-backbone protein mimetics. *Chem. Sci.* **2022**, *13*, 11798–11806.
- (7) Wang, D.; Tiwary, P. Augmenting Human Expertise in Weighted Ensemble Simulations through Deep Learning based Information Bottleneck. *J. Chem. Theory Comput.* **2024**, 10371.
- (8) Leung, J. M. G.; Frazee, N. C.; Brace, A.; Bogetti, A. T.; Ramanathan, A.; Chong, L. T. Unsupervised Learning of Progress Coordinates during Weighted Ensemble Simulations: Application to NTL9 Protein Folding. *J. Chem. Theory Comput.* **2025**, *21*, 3691–3699.
- (9) Saglam, A. S.; Chong, L. T. Protein-protein binding pathways and calculations of rate constants using fully-continuous, explicit-solvent simulations. *Chem. Sci.* **2019**, *10*, 2360–2372.
- (10) Zwier, M. C.; Pratt, A. J.; Adelman, J. L.; Kaus, J. W.; Zuckerman, D. M.; Chong, L. T. Efficient Atomistic Simulation of Pathways and Calculation of Rate Constants for a Protein-Peptide Binding Process: Application to the MDM2 Protein and an Intrinsically Disordered p53 Peptide. *J. Phys. Chem. Lett.* **2016**, *7*, 3440–3445.
- (11) Silvestrini, M. L.; Solazzo, R.; Boral, S.; Cocco, M. J.; Closson, J. D.; Masetti, M.; Gardner, K. H.; Chong, L. T. Gating residues govern ligand unbinding kinetics from the buried cavity in HIF-2 α PAS-B. *Protein Sci.* **2024**, *33*, No. e5198.
- (12) Lotz, S. D.; Dickson, A. Unbiased Molecular Dynamics of 11 min Timescale Drug Unbinding Reveals Transition State Stabilizing Interactions. *J. Am. Chem. Soc.* **2018**, *140*, 618–628.
- (13) Bogetti, A. T.; Zwier, M. C.; Chong, L. T. Revisiting Textbook Azide-Clock Reactions: A “Propeller-Crawling” Mechanism Explains Differences in Rates. *J. Am. Chem. Soc.* **2024**, *146*, 12828–12835.
- (14) Donyapour, N.; Roussey, N. M.; Dickson, A. REVO: Resampling of ensembles by variation optimization. *J. Chem. Phys.* **2019**, *150*, 244112.
- (15) Yang, D. T.; Goldberg, A. M.; Chong, L. T. Rare-Event Sampling using a Reinforcement Learning-Based Weighted Ensemble Method. *bioRxiv*, 2024, DOI: [10.1101/2024.10.09.617475](https://doi.org/10.1101/2024.10.09.617475).
- (16) Zhang, B. W.; Jasnow, D.; Zuckerman, D. M. The “weighted ensemble” path sampling method is statistically exact for a broad class of stochastic processes and binning procedures. *J. Chem. Phys.* **2010**, *132*, No. 054107.

- (17) Nunes-Alves, A.; Kokh, D. B.; Wade, R. C. Ligand unbinding mechanisms and kinetics for T4 lysozyme mutants from tauRAMD simulations. *Curr. Res. Struct. Biol.* **2021**, *3*, 106–111.
- (18) Mollica, L.; Decherchi, S.; Zia, S. R.; Gaspari, R.; Cavalli, A.; Rocchia, W. Kinetics of protein-ligand unbinding via smoothed potential molecular dynamics simulations. *Sci. Rep.* **2015**, *5*, 11539.
- (19) Wolf, S.; Amaral, M.; Lowinski, M.; Vallee, F.; Musil, D.; Guldenhaupt, J.; Dreyer, M. K.; Bomke, J.; Frech, M.; Schlitter, J.; Gerwert, K. Estimation of Protein-Ligand Unbinding Kinetics Using Non-Equilibrium Targeted Molecular Dynamics Simulations. *J. Chem. Inf. Model.* **2019**, *59*, S135–S147.
- (20) Park, S.; Schulten, K. Calculating potentials of mean force from steered molecular dynamics simulations. *J. Chem. Phys.* **2004**, *120*, 5946–5961.
- (21) Tiwary, P.; Limongelli, V.; Salvalaglio, M.; Parrinello, M. Kinetics of protein-ligand unbinding: Predicting pathways, rates, and rate-limiting steps. *Proc. Natl. Acad. Sci. U. S. A.* **2015**, *112*, E386–E391.
- (22) Wang, J.; Miao, Y. Ligand Gaussian Accelerated Molecular Dynamics 3 (LiGaMD3): Improved Calculations of Binding Thermodynamics and Kinetics of Both Small Molecules and Flexible Peptides. *J. Chem. Theory Comput.* **2024**, *20*, 5829–5841.
- (23) Ray, D.; Gokey, T.; Mobley, D. L.; Andricioaei, I. Kinetics and free energy of ligand dissociation using weighted ensemble milestoning. *J. Chem. Phys.* **2020**, *153*, 154117.
- (24) Votapka, L. W.; Ojha, A. A.; Asada, N.; Amaro, R. E. Prediction of Threonine-Tyrosine Kinase Receptor-Ligand Unbinding Kinetics with Multiscale Milestoning and Metadynamics. *J. Phys. Chem. Lett.* **2024**, *15*, 10473–10478.
- (25) Torrillo, P. A.; Bogetti, A. T.; Chong, L. T. A Minimal, Adaptive Binning Scheme for Weighted Ensemble Simulations. *J. Phys. Chem. A* **2021**, *125*, 1642–1649.
- (26) Turner, J.; Anderson, R.; Guo, J.; Beraud, C.; Fletterick, R.; Sakowicz, R. Crystal structure of the mitotic spindle kinesin Eg5 reveals a novel conformation of the neck-linker. *J. Biol. Chem.* **2001**, *276*, 25496–25502.
- (27) Brier, S.; Lemaire, D.; DeBonis, S.; Forest, E.; Kozielski, F. Molecular dissection of the inhibitor binding pocket of mitotic kinesin Eg5 reveals mutants that confer resistance to antimitotic agents. *J. Mol. Biol.* **2006**, *360*, 360–376.
- (28) Cochran, J. C.; Gilbert, S. P. ATPase mechanism of Eg5 in the absence of microtubules: insight into microtubule activation and allosteric inhibition by monastrol. *Biochemistry* **2005**, *44*, 16633–16648.
- (29) Lu, H.; Tonge, P. J. Drug-target residence time: critical information for lead optimization. *Curr. Opin. Chem. Biol.* **2010**, *14*, 467–474.
- (30) Zhang, S.; Thompson, J. P.; Xia, J.; Bogetti, A. T.; York, F.; Skillman, A. G.; Chong, L. T.; LeBard, D. N. Mechanistic Insights into Passive Membrane Permeability of Drug-like Molecules from a Weighted Ensemble of Trajectories. *J. Chem. Inf. Model.* **2022**, *62*, 1891–1904.
- (31) van Erp, T. S.; Moroni, D.; Bolhuis, P. G. A novel path sampling method for the calculation of rate constants. *J. Chem. Phys.* **2003**, *118*, 7762–7774.
- (32) Ray, D.; Stone, S. E.; Andricioaei, I. Markovian Weighted Ensemble Milestoning (M-WEM): Long-Time Kinetics from Short Trajectories. *J. Chem. Theory Comput.* **2022**, *18*, 79–95.
- (33) Case, D. A.; Aktulga, H. M.; Belfon, K.; Cerutti, D. S.; Cisneros, G. A.; Cruzeiro, V. W. D.; Forouzes, N.; Giese, T. J.; Gotz, A. W.; Gohlke, H.; et al. AmberTools. *J. Chem. Inf. Model.* **2023**, *63*, 6183–6191.
- (34) Tian, C.; Kasavajhala, K.; Belfon, K. A. A.; Raguette, L.; Huang, H.; Miguez, A. N.; Bickel, J.; Wang, Y.; Pincay, J.; Wu, Q.; Simmerling, C. ff19SB: Amino-Acid-Specific Protein Backbone Parameters Trained against Quantum Mechanics Energy Surfaces in Solution. *J. Chem. Theory Comput.* **2020**, *16*, 528–552.
- (35) Izadi, S.; Anandakrishnan, R.; Onufriev, A. V. Building Water Models: A Different Approach. *J. Phys. Chem. Lett.* **2014**, *5*, 3863–3871.
- (36) Li, P.; Song, L. F.; Merz, K. M., Jr. Parameterization of highly charged metal ions using the 12–6-4 LJ-type nonbonded model in explicit water. *J. Phys. Chem. B* **2015**, *119*, 883–895.
- (37) Meagher, K. L.; Redman, L. T.; Carlson, H. A. Development of polyphosphate parameters for use with the AMBER force field. *J. Comput. Chem.* **2003**, *24*, 1016–1025.
- (38) Essmann, U.; Perera, L.; Berkowitz, M. L.; Darden, T.; Lee, H.; Pedersen, L. G. A Smooth Particle Mesh Ewald Method. *J. Chem. Phys.* **1995**, *103*, 8577–8593.
- (39) Russo, J. D.; Zhang, S.; Leung, J. M. G.; Bogetti, A. T.; Thompson, J. P.; DeGrave, A. J.; Torrillo, P. A.; Pratt, A. J.; Wong, K. F.; Xia, J.; et al. WESTPA 2.0: High-Performance Upgrades for Weighted Ensemble Simulations and Analysis of Longer-Timescale Applications. *J. Chem. Theory Comput.* **2022**, *18*, 638–649.
- (40) Plotnikov, D.; Ahn, S.-H. Optimization of the resampling method in the weighted ensemble simulation toolkit with parallelization and analysis (WESTPA). *J. Chem. Phys.* **2024**, *161*, No. 046101.